

Joint Covariance Estimation with Mutual Linear Structure

Ilya Soloveychik and Ami Wiesel,

Rachel and Selim Benin School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel

Abstract—We consider the problem of joint estimation of structured covariance matrices. Assuming the structure is unknown, estimation is achieved using heterogeneous training sets. Namely, given groups of measurements coming from centered populations with different covariances, our aim is to determine the mutual structure of these covariance matrices and estimate them. Supposing that the covariances span a low dimensional affine subspace in the space of symmetric matrices, we develop a new efficient algorithm discovering the structure and using it to improve the estimation. Our technique is based on the application of principal component analysis in the matrix space. We also derive an upper performance bound of the proposed algorithm in the Gaussian scenario and compare it with the Cramer-Rao lower bound. Numerical simulations are presented to illustrate the performance benefits of the proposed method.

Index Terms—Structured covariance estimation, joint covariance estimation.

I. INTRODUCTION

Large scale covariance matrix estimation using a small number of measurements is a fundamental problem in modern multivariate statistics. In such environments the Sample Covariance Matrix (SCM) may provide a poor estimate of the true covariance. The scarce amount of samples is especially a problem in financial data analysis where few stationary monthly observations of numerous stock indexes are used to estimate the joint covariance matrix of the stock returns [1, 2], bioinformatics where clustering of genes is obtained based on gene sequences sampled from a small population [3], computational immunology where correlations among mutations in viral strains are estimated from sampled viral sequences and used as a basis of novel vaccine design [4, 5], psychology where the covariance matrix of multiple psychological traits is estimated from data collected on a group of tested individuals [6], or electrical engineering at large where signal samples extracted from a possibly short time window are used to retrieve parameters of the signal [7].

The most common approach to work around the sample deficiency problem is to introduce prior information, which reduces the number of degrees of freedom in the model and allows accurate estimation with few samples. Prior knowledge on the structure can originate from the physics of the underlying phenomena or from similar datasets, see e.g. [8–10]. The

most widely used approaches here are shrinkage towards a known matrix [11, 12], or assuming the true covariance matrix to be structured [13–21]. We focus on prior knowledge in the form of structural constraints, which are most commonly affine. Probably the most popular of them is the Toeplitz model [8, 13–15, 21–24], closely related to it are circulant matrices [16, 17]. In other settings the number of parameters can be reduced by assuming that the covariance matrix is sparse [18–20]. A popular sparse model is the banded covariance, which is associated with time-varying moving average models [15, 21]. Other examples of structured covariances include factor models [25], permutation invariant models [26], and many others.

An important common feature of the works listed in the previous paragraph is that they consider a single and static environment where the structure of the true parameter matrix, or at least the class of structures, as in the sparse case, is known in advance. Often, this is not the case and techniques are needed to learn the structure from the observations. A typical approach is to consider multiple datasets sharing a similar structure but non homogeneous environments [27–29]. This is, for example, the case in covariance estimation for classification across multiple classes [30]. A related problem addresses tracking a time varying covariance throughout a stream of data [15, 31], where it is assumed that the structure changes at a slower rate than the covariances themselves [32]. Here too, it is natural to divide this stream of data into independent blocks of measurements.

Our goal is to first rigorously state the problem of joint covariance estimation with linear structure, and derive lower performance bounds for the family of unbiased estimators. Secondly, we propose and analyze new algorithms of learning and exploiting the linear structure to improve the covariance estimation. More exactly, given a few groups of measurements having different second moments each, our target is to determine the underlying low dimensional affine subspace containing or approximately containing the covariance of all the groups. The discovered hyperplane is further used to improve the matrix parameter estimation. Most of the previous works considered particular cases of this method, e.g. factor models, entry-wise linear structures like in sparse and banded cases, or specific patterns like in Toeplitz, circulant and other models. Our algorithm treats the SCM of the heterogeneous populations as vectors in the space of matrices and is based on application of the principal component analysis (PCA) to learn their low-dimensional structure. To make the performance

This work was partially supported by the Intel Collaboration Research Institute for Computational Intelligence and Israel Science Foundation.

The results were partially presented at the 40-th IEEE International Conference on Acoustics, Speech, and Signal Processing, April, 19-24, 2015, Brisbane, Australia.

analysis a tractable problem, we assume the data is Gaussian in the corresponding sections, however noting that the algorithm itself provides a good estimation in a much wider class of distributions.

The paper is organized as following. First, we introduce the notations, state the problems and provide examples of affine structures motivating the work. Then we derive the lower performance bound for a class of unbiased estimators. Next we propose our algorithm and provide its upper performance bound. We conclude by numerical simulations demonstrating the performance advantages of the proposed technique and supporting our theoretical results. The Appendix section contains all the proofs and auxiliary statements.

A. Notations

Denote by $\mathcal{S}(p)$ the $l = \frac{p(p+1)}{2}$ dimensional linear space of $p \times p$ symmetric real matrices. \mathbf{I}_d stands for the $d \times d$ identity matrix. For a matrix \mathbf{M} its Moore-Penrose generalized inverse is denoted by \mathbf{M}^\dagger . For any two matrices \mathbf{M} and \mathbf{P} we denote by $\mathbf{M} \otimes \mathbf{P}$ their tensor (Kronecker) product. $\sigma_1(\mathbf{M}) \geq \dots \geq \sigma_r(\mathbf{M}) \geq 0$ stand for the singular values of a rectangular matrix \mathbf{M} of rank not greater than r . If $\mathbf{M} \in \mathcal{S}(p)$, we denote its eigenvalues by $\lambda_1(\mathbf{M}) \geq \dots \lambda_p(\mathbf{M})$. $\|\cdot\|_F$ denotes the Frobenius, and $\|\cdot\|_2$ - the spectral norms of matrices. For any symmetric matrix \mathbf{S} , $\mathbf{s} = \text{vech}(\mathbf{S})$ is a vector obtained by stacking the columns of the lower triangular part of \mathbf{S} into a single column. In addition, given an l dimensional column vector \mathbf{m} we denote by $\text{mat}(\mathbf{m})$ the inverse operator constructing a $p \times p$ symmetric matrix such that $\text{vech}(\text{mat}(\mathbf{m})) = \mathbf{m}$. Due to this natural linear bijection below we often consider subsets of $\mathcal{S}(p)$ as subsets of \mathbb{R}^l . In addition, let $\text{vec}(\mathbf{S})$ be a p^2 dimensional vector obtained by stacking the columns of \mathbf{S} , and denote by \mathcal{I} its indices corresponding to the related elements of $\text{vech}(\mathbf{S})$.

ACKNOWLEDGMENT

The authors are grateful to Raj Rao Nadakuditi for his useful remarks and important suggestions, which helped us to significantly improve the paper.

II. PROBLEM FORMULATION AND EXAMPLES

Consider the heterogeneous measurements model, namely, assume we are given $K \geq l = \frac{p(p+1)}{2}$ groups of real p dimensional normal random vectors

$$\mathbf{x}_1^1, \dots, \mathbf{x}_1^n, \dots, \mathbf{x}_K^1, \dots, \mathbf{x}_K^n, \dots, \mathbf{x}_K^1, \dots, \mathbf{x}_K^n \quad (1)$$

with n i.i.d. (independent and identically distributed) samples inside each group and covariances

$$\mathbf{Q}_k = \mathbb{E}[\mathbf{x}_k^1 \mathbf{x}_k^{1T}], \quad k = 1, \dots, K.$$

We assume that

$$\mathbf{Q}_k \in \mathcal{P}(p), \quad k = 1, \dots, K, \quad (2)$$

belong to an r dimensional affine subspace of $\mathcal{S}(p)$. Our main goal is to estimate this subspace and use it to improve the covariance matrices estimation. As a matter of application we

will assume that r is known or not known in advance. In the latter case we will also explain how to determine it from the data.

Let us list the most common affine covariance constraints naturally appearing in typical signal processing applications.

- **Diagonal:** The simplest example of a structure is given by diagonal matrices. The covariance matrix is diagonal when the noise variates are independent, or can be assumed so with great precision. In this case the low dimensional subspace containing the structured matrices is $r = p$ dimensional.
- **Banded:** In a similar manner it is often reasonable to assume non-neighboring variates of the sampled vectors to be uncorrelated. Claiming that i -th element of the random vector is uncorrelated with the h -th if $|i - h| > b$ leads to b -banded covariance structure. The subspace of symmetric b -banded matrices constitutes an $r = \frac{(2p-b)(b+1)}{2}$ dimensional subspace inside $\mathcal{S}(p)$. Banded covariance matrices often naturally appear in time-varying moving average models or in spatially distributed networks, [14, 21–23].
- **Circulant:** The next common type of structured covariance matrices are symmetric circulant matrices, defined as

$$\mathbf{M} = \begin{pmatrix} m_1 & m_2 & m_3 & \dots & m_p \\ m_p & m_1 & m_2 & \dots & m_{p-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ m_2 & m_3 & m_4 & \dots & m_1 \end{pmatrix}, \quad (3)$$

with the natural symmetry conditions such as $m_p = m_2$, etc. Symmetric circulant matrices belong to an $r = p/2$ dimensional subspace if p is even and $(p+1)/2$ if it is odd. Such matrices are typically used as approximations to Toeplitz matrices [9] which are associated with signals that obey periodic stochastic properties for example the yearly variation of temperature in a particular location. A special case of such processes are the classical stationary processes, which are ubiquitous in engineering, [16, 17].

- **Toeplitz:** A natural generalization of circulant are Toeplitz matrices. The covariance matrix appears to possess Toeplitz structure whenever the correlation between the i -th and the h -th components depend only on the the difference $|i - h|$. The dimension of a subspace of Toeplitz symmetric matrices is $r = p$. The two classical models for spectrum estimation utilizing the Toeplitz structures are the moving average (MA) and the autoregressive (AR) processes, [14, 15]. Interestingly, the finite MA(a) process can be easily shown to be equivalent to a -banded Toeplitz covariance model, having the dimension $r = a + 1$.
- **Proper Complex:** Many physical processes can be conveniently described in terms of complex signals. For example, the most frequently appearing model of complex Gaussian noise is the circularly symmetric one [33]. Such noise is completely characterized by its mean and rotation invariant hermitian covariance matrix \mathbf{Q}^C . Denote centered proper complex distributions as

$$\mathbf{x} \sim \mathcal{CN}(\mathbf{0}, \mathbf{Q}^C).$$

The real representation of the covariance reads as

$$\mathbf{Q}^R = \frac{1}{2} \begin{pmatrix} \text{Re}(\mathbf{Q}^C) & -\text{Im}(\mathbf{Q}^C) \\ \text{Im}(\mathbf{Q}^C) & \text{Re}(\mathbf{Q}^C) \end{pmatrix}. \quad (4)$$

We see that \mathbf{Q}^R possesses a simple linear $r = p^2/4$ dimensional structure, where p is the dimension of \mathbf{Q}^R , which is always even. An analogous reasoning applies to proper quaternion covariances [34].

In the following it will be convenient to use a single matrix notation for the multiple linearly structured matrices. Set

$$\mathbf{q}_k = \text{vech}(\mathbf{Q}_k), k = 1, \dots, K, \quad (5)$$

$$\mathbf{Y} = [\mathbf{q}_1, \dots, \mathbf{q}_K], \quad (6)$$

Using these notation, the prior subspace knowledge is equivalent to a low-rank constraint

$$\mathbf{Y} = \mathbf{U}\mathbf{Z}, \quad (7)$$

where $\mathbf{U} \in \mathbb{R}^{l \times r}$ and $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_K] \in \mathbb{R}^{r \times K}$. Essentially our problem reduces to estimation of \mathbf{Y} given K groups of i.i.d. measurements $\{\mathbf{x}_k^i\}_{k,i=1}^{K,n}$ coming from the corresponding centered distributions with covariances $\mathbf{Q}_k, k = 1, \dots, K$.

III. LOWER PERFORMANCE BOUNDS

Before addressing possible solutions for the above joint covariance estimation problem, it is instructive to examine the inherent performance bounds. In this section we assume the rank r is known and the samples are normally distributed. To obtain such performance bounds we use the Cramer-Rao Bound (CRB) to lower bound the Mean Squared Error (MSE) of any unbiased estimator $\hat{\mathbf{Y}}$ of \mathbf{Y} , defined as

$$\text{MSE} = \mathbb{E} \left[\left\| \hat{\mathbf{Y}} - \mathbf{Y} \right\|_F^2 \right]. \quad (8)$$

The MSE is bounded from below by the trace of the corresponding CRB matrix. To compute the CRB, for each i we stack the measurements \mathbf{x}_k^i from (1) into a single vector

$$\mathbf{x}^i = \begin{pmatrix} \mathbf{x}_1^i \\ \vdots \\ \mathbf{x}_K^i \end{pmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{\text{ext}}), \quad i = 1, \dots, n, \quad (9)$$

where the extended covariance is given by

$$\begin{aligned} \mathbf{Q}_{\text{ext}}(\mathbf{U}, \mathbf{Z}) &= \text{diag} \{ \mathbf{Q}_1, \dots, \mathbf{Q}_K \} \\ &= \text{diag} \{ \text{mat}(\mathbf{U}\mathbf{z}_1), \dots, \text{mat}(\mathbf{U}\mathbf{z}_K) \}. \end{aligned} \quad (10)$$

Here the operator $\text{diag} \{ \mathbf{Q}_1, \dots, \mathbf{Q}_K \}$ returns a block-diagonal matrix of size $pK \times pK$ with \mathbf{Q}_k -s as its diagonal blocks. The Jacobian matrix of (9) parametrized as in (10) reads as

$$\begin{aligned} \mathbf{J} &= \frac{\partial \mathbf{Q}_{\text{ext}}}{\partial (\mathbf{U}, \mathbf{Z})} = \begin{pmatrix} \frac{\partial \mathbf{q}_1}{\partial \mathbf{U}} & \frac{\partial \mathbf{q}_1}{\partial \mathbf{z}_1} & 0 & \dots & 0 \\ \frac{\partial \mathbf{q}_2}{\partial \mathbf{U}} & 0 & \frac{\partial \mathbf{q}_2}{\partial \mathbf{z}_2} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{q}_K}{\partial \mathbf{U}} & 0 & 0 & \dots & \frac{\partial \mathbf{q}_K}{\partial \mathbf{z}_K} \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{z}_1^T \otimes \mathbf{I}_l & \mathbf{U} & 0 & \dots & 0 \\ \mathbf{z}_2^T \otimes \mathbf{I}_l & 0 & \mathbf{U} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{z}_K^T \otimes \mathbf{I}_l & 0 & 0 & \dots & \mathbf{U} \end{pmatrix} \in \mathbb{R}^{lK \times (lr + Kr)}, \end{aligned} \quad (11)$$

where we have used the following notation:

$$\frac{\partial \mathbf{q}_k}{\partial \mathbf{U}} = \left[\frac{\partial \mathbf{q}_k}{\partial \mathbf{u}_1} \quad \frac{\partial \mathbf{q}_k}{\partial \mathbf{u}_2} \quad \dots \quad \frac{\partial \mathbf{q}_k}{\partial \mathbf{u}_r} \right], \quad (12)$$

and the formulas

$$\frac{\partial \mathbf{q}_k}{\partial \mathbf{u}_j} = \frac{\partial \mathbf{U}\mathbf{z}_k}{\partial \mathbf{u}_j} = z_k^j \mathbf{I}_l, \quad \frac{\partial \mathbf{q}_k}{\partial \mathbf{z}_k} = \frac{\partial \mathbf{U}\mathbf{z}_k}{\partial \mathbf{z}_k} = \mathbf{U}. \quad (13)$$

Lemma 1.

$$\text{rank}(\mathbf{J}) = lr + Kr - r^2 \quad (14)$$

Proof. The proof can be found in Appendix A. \square

This lemma implies that \mathbf{J} is rank deficient, as

$$\text{rank}(\mathbf{J}) = lr + Kr - r^2 \leq \min[lK, lr + Kr], \quad (15)$$

reflecting the fact that the parametrization of \mathbf{Q}_{ext} or \mathbf{Y} by the pair (\mathbf{U}, \mathbf{Z}) is unidentifiable. Indeed for any invertible matrix \mathbf{A} , the pair $(\mathbf{U}\mathbf{A}, \mathbf{A}^{-1}\mathbf{Z})$ fits as good. Due to this ambiguity the matrix $\mathbf{FIM}(\mathbf{U}, \mathbf{Z})$ is singular and in order to compute the CRB we use the Moore-Penrose pseudo-inverse of $\mathbf{FIM}(\mathbf{U}, \mathbf{Z})$ instead of inverse, as justified by [35]. Given n i.i.d. samples $\mathbf{x}^i, i = 1, \dots, n$, we obtain

$$\text{CRB} = \frac{1}{n} \mathbf{J} \mathbf{FIM}(\mathbf{U}, \mathbf{Z})^\dagger \mathbf{J}^T. \quad (16)$$

For the Gaussian population the matrix $\mathbf{FIM}(\mathbf{U}, \mathbf{Z})$ is given by

$$\mathbf{FIM}(\mathbf{U}, \mathbf{Z}) = \frac{1}{2} \mathbf{J}^T \text{diag} \left\{ [\mathbf{Q}_k^{-1} \otimes \mathbf{Q}_k^{-1}]_{\mathcal{I}, \mathcal{I}} \right\} \mathbf{J}, \quad (17)$$

where $[\mathbf{M}]_{\mathcal{I}, \mathcal{I}}$ is the square submatrix of \mathbf{M} corresponding to the subset of indices from \mathcal{I} . The bound on the MSE is therefore given by

$$\begin{aligned} \text{MSE} &\geq \text{Tr}(\text{CRB}) = \frac{1}{n} \text{Tr}(\mathbf{FIM}(\mathbf{U}, \mathbf{Z})^\dagger \mathbf{J}^T \mathbf{J}) \\ &= \frac{2}{n} \text{Tr} \left(\left[\mathbf{J}^T \text{diag} \left\{ [\mathbf{Q}_k^{-1} \otimes \mathbf{Q}_k^{-1}]_{\mathcal{I}, \mathcal{I}} \right\} \mathbf{J} \right]^\dagger \mathbf{J}^T \mathbf{J} \right). \end{aligned} \quad (18)$$

Denote

$$\underline{\lambda} = \min_k \lambda_p(\mathbf{Q}_k), \quad \bar{\lambda} = \max_k \lambda_1(\mathbf{Q}_k). \quad (19)$$

To get more insight on (18) we bound it from below

$$\begin{aligned} \text{MSE} &\geq \frac{2\underline{\lambda}^2}{n} \text{Tr} \left([\mathbf{J}^T \mathbf{J}]^\dagger \mathbf{J}^T \mathbf{J} \right) \\ &= \frac{2\underline{\lambda}^2}{n} \text{rank}(\mathbf{J}) = \frac{2\underline{\lambda}^2}{n} (lr + Kr - r^2). \end{aligned} \quad (20)$$

The dependence on the model parameters here is similar to that obtained by [36] for the problem of low-rank matrix reconstruction. An important quantity is the marginal MSE per one matrix \mathbf{Q}_k , which is proportional to

$$\frac{\text{MSE}}{K} \sim \frac{lr - r^2}{Kn} + \frac{r}{n}. \quad (21)$$

IV. TSVD ALGORITHM

A. The Basic Algorithm

In this section we present our algorithm for recovery of the true underlying covariances $\mathbf{Q}_1, \dots, \mathbf{Q}_K$. We make use of the representation (7) of \mathbf{Y} with a small change in notation consisting in separation of the columns mean

$$\mathbf{Y} = \mathbf{U}\mathbf{Z} + \mathbf{u}_0 \cdot [1, \dots, 1]. \quad (22)$$

This is done in order to improve the performance of the proposed PCA-based algorithm. Consider the SCM of the k -th group of measurements

$$\mathbf{S}_k = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^k \mathbf{x}_i^{kT}, \quad (23)$$

and let

$$\mathbf{s}_k = \text{vech}(\mathbf{S}_k). \quad (24)$$

Denote by

$$\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_K], \quad (25)$$

the measurement matrix, and compute the average of the columns

$$\hat{\mathbf{u}}_0 = \frac{1}{K} \sum_{k=1}^K \mathbf{s}_k. \quad (26)$$

Consider the matrix

$$\mathbf{S}' = [\mathbf{s}_1 - \hat{\mathbf{u}}_0, \dots, \mathbf{s}_K - \hat{\mathbf{u}}_0], \quad (27)$$

The SVD of \mathbf{S}' reads as

$$\mathbf{S}' = \hat{\mathbf{U}} \begin{pmatrix} \hat{\Sigma} & 0 \\ 0 & \hat{\Sigma}_n \end{pmatrix} \hat{\mathbf{W}}^T = [\hat{\mathbf{U}}_1 \hat{\mathbf{U}}_2] \begin{pmatrix} \hat{\Sigma} & 0 \\ 0 & \hat{\Sigma}_n \end{pmatrix} [\hat{\mathbf{W}}_1 \hat{\mathbf{W}}_2]^T, \quad (28)$$

where the singular values are sorted in the decreasing order and $\hat{\Sigma} \in \mathbb{R}^{r \times r}$. We propose to use the matrix

$$\hat{\mathbf{Y}}' = \hat{\mathbf{U}}_1 \hat{\Sigma} \hat{\mathbf{W}}_1^T, \quad (29)$$

as an estimator of

$$\mathbf{Y}' = \mathbf{U}\mathbf{Z} = [\mathbf{U}_1 \mathbf{U}_2] \begin{pmatrix} \Sigma & 0 \\ 0 & 0 \end{pmatrix} [\mathbf{W}_1 \mathbf{W}_2]^T. \quad (30)$$

This approach is based on Eckart-Young theorem and we refer to it as Truncated SVD (TSVD) method [37]. Finally, for the estimator of \mathbf{Y} we have

$$\hat{\mathbf{Y}} = \hat{\mathbf{Y}}' + \hat{\mathbf{u}}_0 \cdot [1, \dots, 1]. \quad (31)$$

B. How to Choose the Rank?

In real world settings the true rank r of the structure subspace is rarely known in advance and one needs to estimate it from the data before applying the TSVD technique. It is instructive to think about rank estimation, followed by TSVD, simply as thresholding of the data singular values. A large variety of thresholding techniques exist, e.g. hard thresholding, see [38] and references therein. Unfortunately, almost none of them can be applied in our scenario due to two main reasons. First, most of them require the noise to be independent of the signal, which is not the case in our setting. Second, most of

the thresholding approaches rely on prior knowledge about the power or spectral characteristics of the noise, which are known a priori, measured from the secondary data or can be somehow estimated from the samples. In our problem prior information is unavailable and such estimations can not be performed. Instead, we propose a different approach utilizing the fact that the noisy measurements come from Wishart populations.

Consider the expected signal power of \mathbf{S}_k

$$\mathbb{E} \|\mathbf{S}_k\|_F^2 = \frac{n+1}{n} \|\mathbf{Q}_k\|_F^2 + \frac{1}{n} (\text{Tr}(\mathbf{Q}_k))^2. \quad (32)$$

For n large enough

$$\mathbb{E} \|\mathbf{S}_k\|_F^2 \approx \|\mathbf{Q}_k\|_F^2 \left(1 + \frac{c}{\rho(\mathbf{Q}_k)} \right), \quad (33)$$

where $c = \frac{p}{n}$ and

$$\rho(\mathbf{Q}_k) = \frac{p \|\mathbf{Q}_k\|_F^2}{(\text{Tr}(\mathbf{Q}_k))^2} = \cos^{-2}(\angle \mathbf{I}, \mathbf{Q}_k) \quad (34)$$

is the sphericity coefficient, [39], measuring how close is \mathbf{Q}_k to the identity matrix. Now (33) implies that the unknown squared norm of the true covariance $\|\mathbf{Q}_k\|_F^2$ is given by

$$\|\mathbf{Q}_k\|_F^2 = \mathbb{E} \|\mathbf{S}_k\|_F^2 \left(1 + \frac{c}{\rho(\mathbf{Q}_k)} \right)^{-1}. \quad (35)$$

We use the following result to estimate the unknown $\rho(\mathbf{Q}_k)$.

Lemma 2. [39] *Let p and n both tend to infinity in such a way that $\frac{p}{n} \rightarrow c$, then*

$$\rho(\mathbf{S}_k) \rightarrow \rho(\mathbf{Q}_k) + c. \quad (36)$$

Based on this lemma, use $\hat{\rho}(\mathbf{Q}_k) = \rho(\mathbf{S}_k) - c$ as an estimate of $\rho(\mathbf{Q}_k)$ and $\|\mathbf{S}_k\|$ as an estimate of $\mathbb{E} \|\mathbf{S}_k\|_F^2$ to obtain

$$\|\mathbf{Q}_k\|_F^2 \approx \|\mathbf{S}_k\|_F^2 \left(1 - \frac{c}{\rho(\mathbf{S}_k)} \right) = \|\mathbf{S}_k\|_F^2 - \frac{(\text{Tr}(\mathbf{S}_k))^2}{n}. \quad (37)$$

The ratio of the desired signal's power $\|\mathbf{Y}\|_F^2$ to the power of measurements, $\mathbb{E} \|\mathbf{S}\|_F^2$ can now be estimated by

$$\alpha(\mathbf{S}_1, \dots, \mathbf{S}_K) = \frac{\|\mathbf{Y}\|_F^2}{\|\mathbf{S}\|_F^2} = 1 - \frac{1}{n} \frac{\sum_{k=1}^K (\text{Tr}(\mathbf{S}_k))^2}{\sum_{k=1}^K \|\mathbf{S}_k\|_F^2}. \quad (38)$$

This derivation suggests a simple rule of thumb for thresholding the spectrum. As an estimate of the rank \hat{r} we take the number of largest singular values of \mathbf{S} carrying the fraction α of the signal's energy,

$$\hat{r} = \underset{L}{\text{argmin}} \left| \sum_{l=1}^L \sigma_l^2(\mathbf{S}) - \alpha \|\mathbf{S}\|_F^2 \right|. \quad (39)$$

Below we test this method of rank recovery in comparison to different approaches by numerical simulations.

C. TSVD Upper Performance Bound

In this section we provide the performance analysis of the proposed TSVD algorithm in the Gaussian scenario, assuming r and \mathbf{u}_0 are known. We make the following natural assumption.

Assumption 1. *The condition number of the true matrix parameter \mathbf{Y} is bounded from above by*

$$\frac{\sigma_1(\mathbf{Y})}{\sigma_r(\mathbf{Y})} \leq \kappa \sqrt{K}, \quad (40)$$

and the smallest singular value of \mathbf{Y} is bounded from below by

$$\sigma_r(\mathbf{Y}) \geq \epsilon, \quad (41)$$

where κ and ϵ do not depend on K, p, n or r .

In addition, without loss of generality let $\mathbf{u}_0 = 0$.

Theorem 1. *Under Assumption 1 with probability at least $1 - 2e^{-cp}$*

$$\|\hat{\mathbf{Y}} - \mathbf{Y}\|_F^2 \leq C_1 \bar{\lambda}^2 (\kappa + \bar{\lambda})^2 \frac{r(K + C_2 l)}{n}, \quad (42)$$

where c, C_1 and C_2 do not depend on K, p, n or r .

Proof. Denote by \mathbf{R} the zero mean noise matrix

$$\mathbf{R} = \mathbf{S} - \mathbf{Y}. \quad (43)$$

The proof consists of two stages. First, using Weyl's and Wedin's Theorems we establish a deterministic bound on $\|\hat{\mathbf{Y}} - \mathbf{Y}\|_F$ depending on $\|\mathbf{R}\|_2, \sqrt{r}$ and κ . On the second stage we use concentration of measure results to achieve high-probability bound on the norm $\|\mathbf{R}\|_2$, yielding the desired result.

Lemma 3. (Weyl, Theorem 4.11 from [40]) *With the notations of Section IV,*

$$\max_{1 \leq j \leq r} |\sigma_j(\hat{\Sigma}) - \sigma_j(\mathbf{Y})| \leq \|\mathbf{R}\|_2. \quad (44)$$

Lemma 4. (Wedin, Theorem 4.4 from [40]) *With the notations of Section IV,*

$$\sigma_{\min}([\hat{\mathbf{U}}^\perp]^T \mathbf{U}^\perp), \sigma_{\min}(\hat{\mathbf{U}}^T \mathbf{U}) \geq \sqrt{1 - \frac{\|\mathbf{R}\|_2^2}{\sigma_r^2(\mathbf{Y})}}. \quad (45)$$

Lemma 5. *When r and \mathbf{u}_0 are known,*

$$\|\hat{\mathbf{Y}} - \mathbf{Y}\|_F \leq \sqrt{r} \|\mathbf{R}\|_2 \left(\sqrt{2} \frac{\sigma_1(\mathbf{Y}) + 2 + \|\mathbf{R}\|_2}{\sigma_r(\mathbf{Y})} + 1 \right). \quad (46)$$

Proof. Use the triangle inequality, Weyl's and Wedin's theo-

rems to get

$$\begin{aligned} \|\hat{\mathbf{Y}} - \mathbf{Y}\|_F &\leq \|\hat{\mathbf{U}}_1 \hat{\Sigma} \hat{\mathbf{W}}_1 - \mathbf{U}_1 \Sigma \mathbf{W}_1\|_F \\ &+ \|\mathbf{U}_1 \hat{\Sigma} \hat{\mathbf{W}}_1 - \mathbf{U}_1 \Sigma \hat{\mathbf{W}}_1\|_F + \|\mathbf{U}_1 \Sigma \hat{\mathbf{W}}_1 - \mathbf{U}_1 \Sigma \mathbf{W}_1\|_F \\ &\leq \|\hat{\Sigma}\|_2 \|\hat{\mathbf{U}}_1 - \mathbf{U}_1\|_F + \|\hat{\Sigma} - \Sigma\|_F + \sigma_1(\mathbf{Y}) \|\hat{\mathbf{W}}_1 - \mathbf{W}_1\|_F \\ &\leq \left(\|\hat{\Sigma}\|_2 + \sigma_1(\mathbf{Y}) \right) \sqrt{2r} \frac{\|\mathbf{R}\|_2}{\sigma_r(\mathbf{Y})} + \|\hat{\Sigma} - \Sigma\|_F \\ &\leq \sqrt{2r} \frac{1 + \sigma_1(\mathbf{Y})}{\sigma_r(\mathbf{Y})} \|\mathbf{R}\|_2 + \|\hat{\Sigma} - \Sigma\|_2 \left(\sqrt{r} + \frac{\sqrt{2r}}{\sigma_r(\mathbf{Y})} \|\mathbf{R}\|_2 \right) \\ &\leq \sqrt{r} \|\mathbf{R}\|_2 \left(\frac{\sqrt{2}}{\sigma_r(\mathbf{Y})} (\sigma_1(\mathbf{Y}) + 1 + \|\mathbf{R}\|_2) + 1 \right). \end{aligned} \quad (47)$$

□

This result quantifies the intuition that the error bound should depend on the intrinsic dimension r of the estimated subspace, rather than on the ambient dimension l .

We now proceed to the second stage of the proof and develop a high-probability bound on the spectral norm of \mathbf{R} in the following

Lemma 6. *with probability at least $1 - 2e^{-cp}$,*

$$\|\mathbf{R}\|_2 \leq \bar{\lambda} \left(\sqrt{\frac{2}{n}} + C_3 \sqrt{\frac{l}{nK}} \right), \quad (48)$$

where C_3 does not depend on the model parameters K, p, n or r .

Proof. The proof can be found in the Appendix B. □

Plug (48) into (46) to obtain that with probability at least $1 - 2e^{-cp}$,

$$\begin{aligned} \|\hat{\mathbf{Y}} - \mathbf{Y}\|_F &\leq \bar{\lambda} \sqrt{\frac{r}{n}} \left(\sqrt{2} + C_3 \sqrt{\frac{l}{K}} \right) \times \\ &\left(\sqrt{2} \frac{\sigma_1(\mathbf{Y})}{\sigma_r(\mathbf{Y})} + \frac{1}{\sigma_r(\mathbf{Y})} + \frac{\bar{\lambda}}{\sigma_r(\mathbf{Y})} \left(\sqrt{\frac{2}{n}} + C_3 \sqrt{\frac{l}{nK}} \right) + 1 \right) \\ &\leq C_4 \bar{\lambda} (\kappa + \bar{\lambda}) \sqrt{\frac{r}{n}} \left(\sqrt{2K} + C_3 \bar{\lambda} \sqrt{l} \right), \end{aligned} \quad (49)$$

where C_4 does not depend on the model parameters and the last inequality is due to Assumption 1. We finally obtain that with probability at least $1 - 2e^{-cp}$,

$$\|\hat{\mathbf{Y}} - \mathbf{Y}\|_F \leq C_1 \bar{\lambda} (\kappa + \bar{\lambda}) \frac{\sqrt{r}(\sqrt{K} + C_2 \bar{\lambda} \sqrt{l})}{\sqrt{n}}, \quad (50)$$

and the statement follows. □

The obtained performance bound (42) suggests that the error is bounded with high probability by a product of a linear combination of K and l with r , rather than the ambient dimension l of the space. Compared to the bound on the MSE, (20), it shows that the proposed TSVD algorithm exhibits near-optimal dependence on the model parameters K, p, n, r .

V. NUMERICAL SIMULATIONS

A. TSVD Performance in Toeplitz Model

For our first experiment we took a Toeplitz setting with $p = 10$, $\mathbf{U}_0 = \mathbf{I}_p$. The true covariances were positive definite matrices generated as $\mathbf{Q}_k = \mathbf{U}_0 + \sum_{j=1}^r z_k^j \mathbf{D}_j / \|\mathbf{D}_j\|_F$, where \mathbf{D}_j has ones on the j -th and $-j$ -th subdiagonals and zeros otherwise, $r = p - 1$ and z_k^j were i.i.d. uniformly distributed over the interval $[-\frac{1}{2}, \frac{1}{2}]$. Figure 1 shows the dependence of the MSE on n when $K = l$. In the unknown r case we took $\hat{\alpha}$ power threshold for our TSVD algorithm, as defined in (39). For comparison we also plot the MSE-s of the SCM and its projection onto the known subspace structure and the true CRB bound given by (18).

Figure 2 compares different thresholding techniques. As the benchmarks we took the Asymptotically Optimal Hard Thresholding (AOHT), proposed by M. Gavish and D.L. Donoho in [38] and the so-called “elbow method”. The AOHT approach consists in hard thresholding the singular values of a noisy low rank matrix \mathbf{X} at a fixed level of

$$\tau(\mathbf{X}) = \omega\left(\frac{l}{K}\right) \sigma_m(\mathbf{X}), \quad (51)$$

independent of the true rank value. Here the function $\omega(\beta)$ is defined in [38] together with the numerical algorithm of its evaluation, and $\sigma_m(\mathbf{X})$ is the median of the singular values of \mathbf{X} . As the authors demonstrate in [38], this approach provides asymptotically optimal hard thresholding under specific assumptions on the noise (in particular, the noise column vectors are assumed i.i.d. white Gaussian). We applied the threshold $\tau(\mathbf{X})$ to the original measurement \mathbf{S} (AOHT-S in the figure), and its centered counterpart \mathbf{S}' (AOHT-S-c in the figure). As the graph suggests, these two approaches give very close results and perform poorly, compared to our α -thresholding, when the number of measurements n in each group is relatively small. The “elbow method” consists in thresholding the spectrum of \mathbf{X} at the level of the largest gap between the consecutive singular values of \mathbf{X} . This intuitive rule follows the well known observation that the eigenvalues corresponding to the signal and noise usually group into clusters separated by a significant gap. This observation also relies on specific properties of the noise which can be violated in our settings. For comparison, we plot the empirical MSEs of the “elbow method” applied to \mathbf{S} (Elbow-S in the figure) and \mathbf{S}' (Elbow-S-c in the figure) correspondingly.

In the third experiment we set $n = 100$ fixed and explored the dependence of the MSE on the number of groups K in the same setting as before. Figure 3 verifies that the marginal MSE depends on K as predicted by formula (21).

B. TSVD Time-tracking

For the second experiment we considered the problem of tracking a time-varying covariance in complex populations. We used the Data Generating Process (DGP) of Patton and Sheppard, [41], which allows for dynamically changing covariances in the spirit of a multivariate GARCH-type model,




Fig. 1. TSVD algorithm performance.




Fig. 2. TSVD algorithm different thresholding methods.

[42, 43]. One of variations of this DGP suggests the following data model:

$$\mathbf{x}_t = \mathbf{H}_t^{1/2} \mathbf{y}_t, t = 1, \dots, Kn, \quad (52)$$

where we assumed the generating data to be proper complex, $\mathbf{y}_t \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$ and defined the hermitian time-varying covariance \mathbf{H}_t to change according to the law

$$\hat{\mathbf{H}}_t = (1 - \beta) \mathbf{H}_{t-1} + \beta \mathbf{M}_t \mathbf{M}_t^H, \quad (53)$$

$$\mathbf{H}_t = \frac{\hat{\mathbf{H}}_t}{\|\hat{\mathbf{H}}_t\|_F}, t = 1, \dots, Kn. \quad (54)$$




Fig. 3. Marginal TSVD algorithm performance, $n = 100$.




Fig. 4. TSVD algorithm learning the low-dimensional subspace with time.

Here \mathbf{M}_t are random $p \times p$ matrices with i.i.d. standard normally distributed entries, \mathbf{H}_0 is arbitrary positive-definite hermitian and $\beta \in [0, 1]$.

The low-dimensional structure appearing in this setting is due to properness of the covariances (see (4)). In order to explore it, the obtained complex data was represented as double-sized real measurements. Each n clock ticks we formed the SCM $\mathbf{S}_{\frac{t}{n}}$ of the last n measurements, where t was the last time count. Then we concatenated the vector $\text{vech}(\mathbf{S}_{\frac{t}{n}})$ to the matrix $\mathbf{Y}_{\frac{t}{n}-1} \in \mathbb{R}^{l \times (\frac{t}{n}-1)}$ of growing size to obtain $\mathbf{Y}_{\frac{t}{n}}$ and applied our TSVD algorithm to it. Thus, our structure knowledge was updated every n ticks, and we expected the error of the covariance estimation to decrease with time. We performed the experiment with $p = 4, K = 100, n = 30, \beta =$

0.01 and used 90%-power threshold to discover the underlying low-dimensional structure. Figure 4 compares the temporal behavior of the MSE-s of the SCM, TSVD applied to it and the projection of SCM on the subspace spanned by proper covariances. The MSE-s were obtained by averaging the squared errors over 10000 iterations.

VI. CONCLUSION

In this paper we consider the problem of joint covariance estimation with linear structure, given heterogeneous measurements. The main challenge in this scenario is twofold. At first, the underlying structure is to be discovered, and then utilized to improve the covariance estimation. We propose a PCA-based algorithm in the space of matrices, solving the both tasks simultaneously. In addition, we analyze the performance bounds of this algorithm and show it is close to being optimal. Finally, we demonstrate its advantages through numerical simulations.

APPENDIX A

Proof of Lemma 1. Let us calculate the column rank of \mathbf{J} . Complete the columns of \mathbf{U} to an orthonormal basis $\hat{\mathbf{U}}$ of \mathbb{R}^l and multiply \mathbf{J} on the left by an invertible square matrix $\mathbf{I}_K \otimes \hat{\mathbf{U}}$ to obtain

$$\hat{\mathbf{J}} = (\mathbf{I}_K \otimes \hat{\mathbf{U}}) \mathbf{J} = \begin{pmatrix} \mathbf{z}_1^T \otimes \hat{\mathbf{U}} & \mathbf{I}_{l \times r} & 0 & \dots & 0 \\ \mathbf{z}_2^T \otimes \hat{\mathbf{U}} & 0 & \mathbf{I}_{l \times r} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{z}_K^T \otimes \hat{\mathbf{U}} & 0 & 0 & \dots & \mathbf{I}_{l \times r} \end{pmatrix}, \quad (55)$$

where $\mathbf{I}_{l \times r}$ denotes the first r columns of \mathbf{I}_l . Let us look at the first l rows of $\hat{\mathbf{J}}$

$$\hat{\mathbf{J}}_1 = (\mathbf{z}_1^T \otimes \hat{\mathbf{U}} \quad \mathbf{I}_{l \times r} \quad 0 \quad \dots \quad 0). \quad (56)$$

By subtracting from the first l columns necessary linear combinations of $\mathbf{I}_{l \times r}$ this matrix can be brought to the form

$$\hat{\mathbf{J}}'_1 = (\mathbf{z}_1^T \otimes \hat{\mathbf{U}}' \quad \mathbf{I}_{l \times r} \quad 0 \quad \dots \quad 0), \quad (57)$$

where the first r rows of $\hat{\mathbf{U}}'$ are zero. Performing the same procedure for all the row submatrices $\hat{\mathbf{J}}_k, k = 1, \dots, K$ yields

$$\hat{\mathbf{J}}' = \begin{pmatrix} \mathbf{z}_1^T \otimes \hat{\mathbf{U}}' & \mathbf{I}_{l \times r} & 0 & \dots & 0 \\ \mathbf{z}_2^T \otimes \hat{\mathbf{U}}' & 0 & \mathbf{I}_{l \times r} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{z}_K^T \otimes \hat{\mathbf{U}}' & 0 & 0 & \dots & \mathbf{I}_{l \times r} \end{pmatrix}, \quad (58)$$

of the same rank as $\hat{\mathbf{J}}$. The first l columns of $\hat{\mathbf{J}}'$ read as

$$\hat{\mathbf{J}}'^1 = \begin{pmatrix} z_{11} \hat{\mathbf{U}}' \\ z_{21} \hat{\mathbf{U}}' \\ \vdots \\ z_{K1} \hat{\mathbf{U}}' \end{pmatrix} = \begin{pmatrix} z_{11} \\ z_{21} \\ \vdots \\ z_{K1} \end{pmatrix} \otimes \hat{\mathbf{U}}', \quad (59)$$

and therefore,

$$\text{rank}(\hat{\mathbf{J}}'^1) = \text{rank}(\hat{\mathbf{U}}') = l - r. \quad (60)$$

Performing the same operation on the column submatrices $\hat{\mathbf{J}}^{tj}$, $j = 1, \dots, r$ gives

$$\text{rank}(\mathbf{J}) = \text{rank}(\hat{\mathbf{J}}') = r(l - r) + Kr = lr + Kr - r^2. \quad (61)$$

□

APPENDIX B

The following definitions and basic lemmas from the theory of Orlicz spaces can be found in such classical references as [44], or in a more convenient form in recent works such as [45].

Definition 1. (Subexponential variable and its Orlicz norm) Let x be a nonnegative random variable such that

$$\mathbb{P}(x \geq \mu + t) \leq \begin{cases} \exp\left(-\frac{t^2}{2\sigma^2}\right), & t \in \left[0, \frac{\sigma^2}{b}\right], \\ \exp\left(-\frac{t}{2b}\right), & t \in \left(\frac{\sigma^2}{b}, \infty\right), \end{cases} \quad (62)$$

for some μ, σ and b , then we say that x is subexponential with parameters (σ^2, b) . The Orlicz norm of x is defined as

$$\|x\|_{\psi_1} = \inf_{\eta > 0} \left\{ \mathbb{E} \exp\left(\frac{x}{\eta}\right) \leq 2 \right\}. \quad (63)$$

The intuition behind the subexponential norm basically says that if $\|x\|_{\psi_1}$ is finite, the $\|x\|_{\psi_1}$ -scaled tails of x are not heavier than that of a standard exponential variable.

Lemma 7. A random variable x is subexponential with parameters (σ^2, b) iff

$$\mathbb{E}(\exp(\lambda x)) \leq \exp\left(\frac{\sigma^2 \lambda^2}{2}\right), \quad \forall |\lambda| < \frac{1}{b}. \quad (64)$$

Corollary 1. The Orlicz norm of a subexponential variable x with parameters (σ^2, b) is bounded by

$$\|x\|_{\psi_1} \leq \max\left(\frac{\sigma}{2\ln(2)}, b\right), \quad (65)$$

Proof. Lemma 7 implies

$$\mathbb{E} \exp\left(\frac{x}{\kappa}\right) \leq \exp\left(\frac{\sigma^2}{2\kappa^2}\right), \quad \forall \kappa \geq b. \quad (66)$$

Bound the right hand side of the last inequality by 2 to get the statement. □

Next we use the following tail bound on a χ^2 variable to calculate the Orlicz norm of its centered version.

Lemma 8. [46] Let $\zeta \sim \chi^2(p)$, then

$$\mathbb{P}(\zeta - p \geq 2\sqrt{pt} + 2t) \leq e^{-t}, \quad (67)$$

$$\mathbb{P}(p - \zeta \geq 2\sqrt{pt}) \leq e^{-t}. \quad (68)$$

Corollary 2. Let $\zeta \sim \chi^2(p)$, then the Orlicz norm of $|\zeta - p|$ is bounded by

$$\|\zeta - p\|_{\psi_1} \leq 2\sqrt{p}. \quad (69)$$

Proof. Follows by a straight forward calculation from Lemma 8 and Corollary 1. □

In the course of proof of Theorem 1 we also use the following tail bound on the spectral norm deviations of a standard Wishart matrix from its mean.

Lemma 9. [47] Let $\mathbf{W} = \frac{1}{n} \sum_{i=1}^n \mathbf{w}_i \mathbf{w}_i^T$, where \mathbf{w}_i are i.i.d. copies of $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I})$, then

$$\mathbb{P}\left(\|\mathbf{W} - \mathbf{I}\|_2 \geq \sqrt{\frac{p}{n}} + t\right) \leq 2e^{-nt^2/2}, \quad (70)$$

Theorem 2 (Theorem 1 from [48]). Let $\mathbf{r}_1, \dots, \mathbf{r}_K \in \mathbb{R}^l$ be independent random vectors (not necessarily identically distributed). Assume that there exists ψ , such that

$$\max_{1 \leq k \leq K} \max_{\mathbf{p} \in \mathcal{S}^{l-1}} \|\langle \mathbf{r}_k, \mathbf{p} \rangle\|_{\psi_1} \leq \psi, \quad (71)$$

where $\mathcal{S}^{l-1} \subset \mathbb{R}^l$ is a unit sphere, and there exists $\mathcal{K} > 1$, such that

$$\mathbb{P}\left(\frac{1}{\sqrt{l}} \max_{1 \leq k \leq K} \|\mathbf{r}_k\| \geq \mathcal{K} \max\left[1, \left(\frac{K}{l}\right)^{1/4}\right]\right) \leq e^{-\sqrt{l}}, \quad (72)$$

then with probability at least $1 - 2e^{-c\sqrt{l}}$

$$\|\mathbf{R}\|_2 \leq \|\mathbf{E}[\mathbf{R}\mathbf{R}^T]\|_2^{1/2} + C(\psi + \mathcal{K})\sqrt{\frac{l}{K}}, \quad (73)$$

where $\mathbf{R} = [\mathbf{r}_1, \dots, \mathbf{r}_K]$ and c, C are universal constants.

Proof of Lemma 6. To make the calculations easier in this proof we assume that the columns of \mathbf{R} ,

$$\mathbf{R} = [\mathbf{r}_1, \dots, \mathbf{r}_K], \quad (74)$$

are constructed as $\mathbf{r}_k = \text{vec}(\mathbf{S}_k - \mathbf{Q}_k)$ and therefore, $\mathbf{R} \in \mathbb{R}^{p^2 \times K}$. This may only affect the numerical constants by the magnitude up to 2. In order to apply Theorem 2 we need to make sure the conditions (71) and (72) hold. For this purpose fix a matrix \mathbf{P} and consider the univariate variable $|\langle \mathbf{r}_k, \mathbf{p} \rangle|$, where $\mathbf{p} = \text{vec}(\mathbf{P})$, for some fixed k . Note that

$$\langle \mathbf{r}_k, \mathbf{p} \rangle = \text{Tr}((\mathbf{S}_k - \mathbf{Q}_k)\mathbf{P}) = \sum_i \mathbf{x}_k^{iT} \mathbf{P} \mathbf{x}_k^i - \text{Tr}(\mathbf{Q}_k \mathbf{P}), \quad (75)$$

and among all the matrices \mathbf{P} of norm one, the Orlicz norm of the right-hand side of (75) is maximized by

$$\mathbf{P}_k = \frac{\mathbf{Q}_k^{-1}}{\|\mathbf{Q}_k^{-1}\|_F}. \quad (76)$$

For this choice of \mathbf{P}_k we obtain

$$|\text{Tr}((\mathbf{S}_k - \mathbf{Q}_k)\mathbf{P}_k)| = \frac{|\frac{1}{n} \sum_i \|\mathbf{x}_k^i\|^2 - p|}{\|\mathbf{Q}_k^{-1}\|_F} \sim \frac{\frac{1}{n} |\zeta - np|}{\|\mathbf{Q}_k^{-1}\|_F}, \quad (77)$$

where $\zeta \sim \chi^2(np)$ and, thus,

$$\max_{\|\mathbf{P}\|_F=1} \|\text{Tr}((\mathbf{S}_k - \mathbf{Q}_k)\mathbf{P})\|_{\psi_1} \leq \frac{2\sqrt{np}}{n \|\mathbf{Q}_k^{-1}\|_F} \leq \frac{2\lambda_1(\mathbf{Q}_k)}{\sqrt{n}}, \quad (78)$$

where we have used the inequality

$$\|\mathbf{Q}_k^{-1}\|_F \geq \frac{\sqrt{p}}{\lambda_1(\mathbf{Q}_k)}. \quad (79)$$

The bound (78) finally yields

$$\max_k \max_{\|\mathbf{P}\|_F=1} \|\text{Tr}((\mathbf{S}_k - \mathbf{Q}_k)\mathbf{P})\|_{\psi_1} \leq \frac{2\bar{\lambda}}{\sqrt{n}}, \quad (80)$$

and therefore the (71) holds true for

$$\psi = \frac{2\bar{\lambda}}{\sqrt{n}}. \quad (81)$$

In order to verify the boundedness condition (72), consider the variable

$$\max_k \frac{\|\mathbf{r}_k\|}{\sqrt{p^2}} = \max_k \frac{\|\mathbf{S}_k - \mathbf{Q}_k\|_F}{p} \leq \frac{\bar{\lambda}}{\sqrt{p}} \max_k \|\mathbf{W}_k - \mathbf{I}\|_2, \quad (82)$$

where $\mathbf{W}_k = \mathbf{Q}_k^{-1/2} \mathbf{S}_k \mathbf{Q}_k^{-1/2}$. Denote

$$\delta = \max \left[1, \left(\frac{K}{p^2} \right)^{1/4} \right] \quad (83)$$

and use Lemma 9 to obtain

$$\begin{aligned} \mathbb{P} \left(\frac{1}{p} \max_k \|\mathbf{r}_k\| \geq \mathcal{K}\delta \right) &\leq \mathbb{P} \left(\frac{\bar{\lambda}}{\sqrt{p}} \max_k \|\mathbf{W}_k - \mathbf{I}\|_2 \geq \mathcal{K}\delta \right) \\ &\leq K \mathbb{P} \left(\|\mathbf{W}_k - \mathbf{I}\|_2 \geq \mathcal{K}\delta \frac{\sqrt{p}}{\bar{\lambda}} \right) \\ &\leq 2K \exp \left(-\frac{n}{2} \left[\mathcal{K}\delta \frac{\sqrt{p}}{\bar{\lambda}} - \sqrt{\frac{p}{n}} \right]^2 \right). \end{aligned} \quad (84)$$

Set

$$\mathcal{K} = \frac{\bar{\lambda}}{\sqrt{n\delta}} (\sqrt{2(1 + \ln 2K/p)} + 1) \quad (85)$$

to get

$$\mathbb{P} \left(\frac{1}{p} \max_k \|\mathbf{r}_k\| \geq \mathcal{K}\delta \right) \leq e^{-p}, \quad (86)$$

thus the conditions of Theorem 1 from [48] are satisfied with the constants ϕ and \mathcal{K} . Let us bound the spectral norm of $\mathbb{E}[\mathbf{R}\mathbf{R}^T]$. For this purpose note that

$$\|\mathbb{E}[\mathbf{R}\mathbf{R}^T]\|_2 \leq \bar{\lambda}^2 \|\mathbb{E}[\mathbf{M}\mathbf{M}^T]\|_2, \quad (87)$$

where

$$\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_K], \quad (88)$$

and $\mathbf{m}_k = \text{vec}(\mathbf{W}_k - \mathbf{I})$, with $\mathbf{W}_k = \frac{1}{n} \sum_{i=1}^n \mathbf{w}_i^k \mathbf{w}_i^{kT}$, and $\mathbf{w}_i^k \sim \mathcal{N}(0, \mathbf{I})$, $i = 1, \dots, n$, $k = 1, \dots, K$. Note that for $\zeta \sim \chi^2(n)$,

$$\mathbb{E}[(\zeta - n)^2] = 2n, \quad (89)$$

and for $\gamma_i, \omega_i \sim \mathcal{N}(0, 1)$, $i = 1, \dots, n$, i.i.d.,

$$\mathbb{E} \left[\left(\sum_{i=1}^n \gamma_i \omega_i \right)^2 \right] = n, \quad (90)$$

to obtain

$$\mathbb{E}[\mathbf{M}\mathbf{M}^T] = \frac{1}{n^2} \begin{pmatrix} 2n\mathbf{I}_p & 0 & 0 \\ 0 & n\mathbf{I}_{l-p} & n\mathbf{I}_{l-p} \\ 0 & n\mathbf{I}_{l-p} & n\mathbf{I}_{l-p} \end{pmatrix} \in \mathbb{R}^{p^2 \times p^2}, \quad (91)$$

where for convenience we have ordered the elements of $\text{vec}(\mathbf{A})$ in such a way that the diagonal of \mathbf{A} goes first, then the upper triangular part and then the lower triangular part. We, therefore, obtain,

$$\|\mathbb{E}[\mathbf{R}\mathbf{R}^T]\|_2 \leq \frac{2\bar{\lambda}^2}{n}. \quad (92)$$

Finally we obtain that with probability at least $1 - 2e^{-cp}$,

$$\begin{aligned} \|\mathbf{R}\|_2 &\leq \bar{\lambda} \sqrt{\frac{2}{n}} + C_2(\phi + \mathcal{K}) \sqrt{\frac{p^2}{K}} \\ &= \bar{\lambda} \sqrt{\frac{2}{n}} + \frac{C_2 \bar{\lambda}}{\sqrt{n}} \left(2 + \frac{\sqrt{2(1 + \ln 2K/p)} + 1}{\delta} \right) \sqrt{\frac{p^2}{K}} \\ &\leq \bar{\lambda} \left(\sqrt{\frac{2}{n}} + C_3 \sqrt{\frac{p^2}{nK}} \right), \end{aligned} \quad (93)$$

where we have used the fact that

$$\frac{\sqrt{2(1 + \ln 2K/p)} + 1}{\delta} = \frac{\sqrt{2(1 + \ln 2K/p)} + 1}{\max \left[1, \left(\frac{K}{p^2} \right)^{1/4} \right]} \quad (94)$$

is bounded. Now replace p^2 back by l , which can at most affect the constants by a factor up to 2, to get the statement. \square

REFERENCES

- [1] L. Laloux, P. Cizeau, M. Potters, and J.-P. Bouchaud, "Random matrix theory and financial correlations," *International Journal of Theoretical and Applied Finance*, vol. 3, no. 03, pp. 391–397, 2000.
- [2] O. Ledoit and M. Wolf, "Improved estimation of the covariance matrix of stock returns with an application to portfolio selection," *Journal of empirical finance*, vol. 10, no. 5, pp. 603–621, 2003.
- [3] J. Schäfer and K. Strimmer, "A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics," *Statistical applications in genetics and molecular biology*, vol. 4, no. 1, 2005.
- [4] V. Dahirel, K. Shekhar, F. Pereyra, T. Miura, M. Artyomov, S. Talsania, T. M. Allen, M. Altfeld, M. Carrington, D. J. Irvine *et al.*, "Coordinate linkage of HIV evolution reveals regions of immunological vulnerability," *Proceedings of the National Academy of Sciences*, pp. 11 530–11 535, 2011.
- [5] A. A. Quadeer, R. H. Louie, K. Shekhar, A. K. Chakraborty, I. M. Hsing, and M. R. McKay, "Statistical linkage of mutations in the non-structural proteins of hepatitis c virus exposes targets for immunogen design," *submitted to Journal of Virology*, 2013.
- [6] J. H. Steiger, "Tests for comparing elements of a correlation matrix," *Psychological Bulletin*, vol. 87, no. 2, p. 245, 1980.
- [7] L. L. Scharf, "Statistical signal processing: Detection, estimation and time-series analysis," 1991.
- [8] D. R. Fuhrmann, "Application of Toeplitz covariance estimation to adaptive beamforming and detection," *IEEE Transactions on Signal Processing*, vol. 39, no. 10, pp. 2194–2198, 1991.
- [9] D. S. Pollock, "Circulant matrices and time-series analysis," *International Journal of Mathematical Education in Science and Technology*, vol. 33, no. 2, pp. 213–230, 2002.
- [10] P. Stoica, P. Babu, and J. Li, "SPICE: A sparse covariance-based estimation method for array processing," *IEEE Transactions on Signal Processing*, vol. 59, no. 2, pp. 629–638, 2011.

- [11] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *Journal of multivariate analysis*, vol. 88, no. 2, pp. 365–411, 2004.
- [12] Y. Chen, A. Wiesel, and A. O. Hero, "Robust shrinkage estimation of high-dimensional covariance matrices," *IEEE Transactions on Signal Processing*, vol. 59, no. 9, pp. 4097–4107, 2011.
- [13] D. L. Snyder, J. A. O'Sullivan, and M. I. Miller, "The use of maximum likelihood estimation for forming images of diffuse radar targets from delay-Doppler data," *IEEE Transactions on Information Theory*, vol. 35, no. 3, pp. 536–548, 1989.
- [14] Y. I. Abramovich, N. K. Spencer, and M. D. Turley, "Time-varying autoregressive (TVAR) models for multiple radar observations," *IEEE Transactions on Signal Processing*, vol. 55, no. 4, pp. 1298–1311, 2007.
- [15] A. Wiesel, O. Bibi, and A. Globerson, "Time varying autoregressive moving average models for covariance estimation," *IEEE Transactions on Signal Processing*, vol. 61, no. 11, pp. 2791–2801, 2013.
- [16] A. Dembo, C. Mallows, and L. Shepp, "Embedding nonnegative definite Toeplitz matrices in nonnegative definite circulant matrices, with application to covariance estimation," *IEEE Transactions on Information Theory*, vol. 35, no. 6, pp. 1206–1212, 1989.
- [17] T. T. Cai, Z. Ren, and H. H. Zhou, "Optimal rates of convergence for estimating Toeplitz covariance matrices," *Probability Theory and Related Fields*, vol. 156, no. 1–2, pp. 101–143, 2013.
- [18] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.
- [19] O. Banerjee, L. El Ghaoui, and A. d'Aspremont, "Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data," *The Journal of Machine Learning Research*, vol. 9, pp. 485–516, 2008.
- [20] T. T. Cai, Z. Ren, and H. H. Zhou, "Optimal rates of convergence for estimating Toeplitz covariance matrices," *Probability Theory and Related Fields*, pp. 1–43, 2012.
- [21] P. J. Bickel and E. Levina, "Regularized estimation of large covariance matrices," *The Annals of Statistics*, pp. 199–227, 2008.
- [22] A. Kavcic and J. M. F. Moura, "Matrices with banded inverses: Inversion algorithms and factorization of Gauss-Markov processes," *IEEE transactions on Information Theory*, vol. 46, no. 4, pp. 1495–1509, 2000.
- [23] A. Asif and J. M. F. Moura, "Block matrices with L-block-banded inverse: inversion algorithms," *IEEE Transactions on Signal Processing*, vol. 53, no. 2, pp. 630–642, 2005.
- [24] W. J. Roberts and Y. Ephraim, "Hidden Markov modeling of speech using Toeplitz covariance matrices," *Speech Communication*, vol. 31, no. 1, pp. 1–14, 2000.
- [25] R. F. Engle, V. K. Ng, and M. Rothschild, "Asset pricing with a factor-ARCH covariance structure: Empirical estimates for treasury bills," *Journal of Econometrics*, vol. 45, no. 1, pp. 213–237, 1990.
- [26] P. Shah and V. Chandrasekaran, "Group symmetry and covariance regularization," *Electronic Journal of Statistics*, vol. 6, pp. 1600–1640, 2012.
- [27] J. Guo, E. Levina, G. Michailidis, and J. Zhu, "Joint estimation of multiple graphical models," *Biometrika*, vol. 98, no. 1, pp. 1–15, 2011.
- [28] O. Besson, S. Bidon, and J.-Y. Tournet, "Covariance matrix estimation with heterogeneous samples," *Signal Processing, IEEE Transactions on*, vol. 56, no. 3, pp. 909–920, 2008.
- [29] S. Bidon, O. Besson, and J.-Y. Tournet, "A Bayesian approach to adaptive detection in nonhomogeneous environments," *Signal Processing, IEEE Transactions on*, vol. 56, no. 1, pp. 205–217, 2008.
- [30] P. Danaher, P. Wang, and D. M. Witten, "The joint graphical lasso for inverse covariance estimation across multiple classes," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 76, no. 2, pp. 373–397, 2014.
- [31] E. Moulines, P. Priouret, F. Roueff *et al.*, "On recursive estimation for time varying autoregressive processes," *The Annals of statistics*, vol. 33, no. 6, pp. 2610–2654, 2005.
- [32] A. Ahmed and E. P. Xing, "Recovering time-varying networks of dependencies in social and biological studies," *Proceedings of the National Academy of Sciences*, vol. 106, no. 29, pp. 11 878–11 883, 2009.
- [33] S. M. Kay, "Fundamentals of statistical signal processing, volume i: Estimation theory (v. 1)," 1993.
- [34] A. Sloin and A. Wiesel, "Proper quaternion gaussian graphical models," 2014.
- [35] Y.-H. Li and P.-C. Yeh, "An interpretation of the Moore-Penrose generalized inverse of a singular Fisher Information Matrix," *IEEE Transactions on Signal Processing*, vol. 60, no. 10, pp. 5532–5536, 2012.
- [36] G. Tang and A. Nehorai, "Lower bounds on the mean-squared error of low-rank matrix reconstruction," *IEEE Transactions on Signal Processing*, vol. 59, no. 10, pp. 4559–4571, 2011.
- [37] G. W. Stewart, "On the early history of the singular value decomposition," *SIAM review*, vol. 35, no. 4, pp. 551–566, 1993.
- [38] D. L. Donoho and M. Gavish, "The optimal hard threshold for singular values is $4/\sqrt{3}$," *arXiv preprint arXiv:1305.5870*, 2013.
- [39] O. Ledoit and M. Wolf, "Some hypothesis tests for the covariance matrix when the dimension is large compared to the sample size," *Annals of Statistics*, pp. 1081–1102, 2002.
- [40] G. W. Stewart and J.-G. Sun, "Matrix perturbation theory," 1990.
- [41] A. J. Patton and K. Sheppard, "Evaluating volatility and correlation forecasts," pp. 801–838, 2009.
- [42] T. Bollerslev, R. F. Engle, and J. M. Wooldridge, "A capital asset pricing model with time-varying covariances," *The Journal of Political Economy*, pp. 116–131, 1988.
- [43] B. Hansson and P. Hordahl, "Testing the conditional CAPM using multivariate GARCH-M," *Applied Financial Economics*, vol. 8, no. 4, pp. 377–388, 1998.

- [44] M. M. Rao and Z. D. Ren, “Theory of Orlicz spaces,” 1991.
- [45] S. van de Geer and J. Lederer, “The Bernstein-Orlicz norm and deviation inequalities,” *Probability Theory and Related Fields*, vol. 157, no. 1-2, pp. 225–250, 2013.
- [46] B. Laurent and P. Massart, “Adaptive estimation of a quadratic functional by model selection,” *Annals of Statistics*, pp. 1302–1338, 2000.
- [47] R. Vershynin, “Introduction to the non-asymptotic analysis of random matrices,” *arXiv preprint arXiv:1011.3027*, 2010.
- [48] R. Adamczak, A. E. Litvak, A. Pajor, and N. Tomczak-Jaegermann, “Sharp bounds on the rate of convergence of the empirical covariance matrix,” *Comptes Rendus Mathematique*, vol. 349, no. 3, pp. 195–200, 2011.